YEARS DIW BERLIN

2110

Discussion Papers

Deutsches Institut für Wirtschaftsforschung 2025

Time-Varying Shock Transmission in Non-Gaussian Structural Vector Autoregressions

Helmut Lütkepohl and Till Strohsal

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

IMPRESSUM

© DIW Berlin, 2025

DIW Berlin
German Institute for Economic Research
Mohrenstr. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
http://www.diw.de

ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
http://www.diw.de/discussionpapers

Discussion Papers of DIW Berlin are indexed in RePEc and Econstor:
http://ideas.repec.org/s/diw/diwwpp.html
https://www.econstor.eu/handle/10419/10

Time-Varying Shock Transmission in Non-Gaussian Structural Vector Autoregressions

Helmut Lütkepohl
DIW Berlin, Mohrenstr. 58, 10117 Berlin, Germany
email: luetkepohlh@gmail.com

Till Strohsal
HWR Berlin and Freie Universität Berlin, Badensche Str. 52, 10825 Berlin, Germany
email: till.strohsal@hwr-berlin.de

February 28, 2025[1]

This paper analyzes possibly time-varying shock transmission in structural vector autoregressive (VAR) models when the reduced-form VAR coefficients are time-invariant and the shocks are identified through non-Gaussianity. To check for possible time-variation in the impulse responses, we propose Wald tests for two situations: (1) homoskedastic and (2) heteroskedastic structural shocks. For the latter case, the challenge is to ensure that the test does not indicate time-varying impulse responses if the changes are due only to changes in the variances of the shocks. To illustrate the usefulness of the tests, they are applied to an empirical model of the crude oil market. They support time-varying shock transmission reflected in impulse response functions that change over time.

Key Words: Structural vector autoregression, independent component analysis, non-Gaussian shocks, structural break tests, heteroskedasticity.

JEL classification: C32

[1]The authors thank Lars Winkelmann for valuable comments on an earlier version of the paper.

1 Introduction

Identification of structural shocks through non-Gaussianity in structural vector autoregressive (VAR) analysis is considered in a number of recent articles (e.g., Lanne and Lütkepohl (2010), Herwartz and Plödt (2016), Herwartz (2018), Hafner, Herwartz and Wang (2025), Hafner and Herwartz (2023), Gouriéroux, Monfort and Renne (2017), Gourieroux and Monfort (2014), Moneta, Entner, Hoyer and Coad (2013), Lanne, Meitz and Saikkonen (2017), Maxand (2020)). While these authors use different approaches for modelling and estimation, the main idea for achieving identification is that a nontrivial linear transformation of a vector of independent non-Gaussian random variables leads to a dependent random vector. Here we refer to a linear transformation as trivial if it just multiplies the independent components by a constant and possibly reorders them. In turn, if u is a K-dimensional non-Gaussian random vector and there exists a linear transformation $u = Bw$ such that the $(K \times 1)$ random vector $w = (w_1, \dots, w_K)'$ has non-Gaussian independent components with unit variances, then the transformation matrix B is unique apart from column sign and column permutations. This result is the basis for identifying shocks in structural VAR analysis via non-Gaussianity.

Finding the transformation matrix B for a given random vector with dependent components from a sample $u_1, \dots, u_T$ from the distribution of u is known as independent component analysis (ICA). Hyvärinen, Karhunen and Oja (2001) review a number of numerical procedures for ICA. As a nonlinear transformation may be needed to extract the independent components, identification in structural VAR analysis based on non-Gaussianity of the data hinges on the existence of a linear transformation that leads to independent components.

An important assumption common in this literature is that the VAR residuals u_t and the shocks w_t have the same distribution across the sample and this feature is then used to estimate the transformation matrix B. Such an assumption is convenient because, if the slope coefficients of the reduced-form VAR process are time-invariant across the sample, it implies time-invariant structural impulse responses. The assumption is not necessary in this framework, however, because non-Gaussianity allows to identify different transformation matrices and impulse responses across the sample if the distribution changes but remains non-Gaussian. For example, there may be heteroskedasticity which implies changes in the variances and, hence, the distributions of the shocks and may imply changes in the transmission of the shocks.

In related work, if the data are heteroskedastic or conditionally heteroskedastic, some authors use this feature for identifying shocks in structural VAR analysis (e.g., Lütkepohl and Netšunajev (2017), Kilian and Lütkepohl (2017, Chapter 14)). In this literature the shocks are also typically chosen as linear transformations of the reduced-form errors but it is just required that the structural shocks are uncorrelated rather than independent. For identifying the structural shocks, some time-invariance of the structural impulse responses is crucial in this literature, however.

In this study we use non-Gaussianity for identification but explicitly allow for time-varying impulse responses despite time-invariant VAR slope coefficients. We propose tests for time-varying shock transmission and explore their properties under different types of changes in the distribution of the VAR residuals. To illustrate the usefulness of the tests for time-varying shock transmission for applied empirical work, we present an example that shows the importance of allowing for distributional heterogeneity. More precisely, we consider a structural VAR model for the crude oil market.

The remainder of this study is organized as follows. The model setup and ICA techniques are presented in the next section. In Section 3, the testing problem is laid out and tests are presented. Their small sample properties are explored in Section 4. The empirical example is discussed in Section 5. Section 6 concludes. Details of the data generating processes (DGPs) used in the small sample simulations and of some bootstrap procedures are provided in Appendices.

2 Non-Gaussian SVAR Models

2.1 Model Setup

Let $u_t = (u_{1t}, \dots, u_{Kt})'$ be a K-dimensional vector of random variables with mean zero and covariance matrix Σ_u, i.e., $u_t \sim (0, \Sigma_u)$. Suppose there exists a nonsingular $(K \times K)$ matrix B such that

$$u_t = Bw_t, \tag{1}$$

where the components of $w_t = (w_{1t}, \dots, w_{Kt})'$ are stochastically independent, have variance one, $w_t \sim (0, I_K)$, and at most one component has a Gaussian distribution. Then B is unique up to column permutations and column sign. In other words, if there exists a matrix B^* and a random vector with independent components $w_t^* \sim (0, I_K)$ such that

$$u_t = B^* w_t^*,$$

then $B^* = BPD$, where P is a permutation matrix and D is a diagonal matrix with diagonal elements ± 1 so that the elements of w_t^* are the same as those of w_t but may have reversed sign and may be ordered differently. In other words, $w_t = PDw_t^*$, where P is a permutation matrix and D is a diagonal matrix with diagonal elements ± 1 (see Comon, Jutten and Herault (1991), Comon (1994), Gouriéroux et al. (2017)).

This result does not mean that for every K-dimensional non-Gaussian stochastic vector u_t with dependent components, there exists a linear transformation that transforms the components to independent random variables. In other words, there may not exist a matrix B such that $B^{-1}u_t$ has independent components. However, if such a matrix exists, then it is unique apart from column permutations and column signs.

The uniqueness (up to column sign and order) of B in equation (1) is used in the structural VAR literature to identify structural shocks. Suppose that

$$y_t = A_1 y_{t-1} + \cdots + A_p y_{t-p} + u_t, \tag{2}$$

where u_t is a zero mean white noise process with covariance matrix Σ_u, i.e., $u_t \sim (0, \Sigma_u)$. Then it is assumed that the vector of structural errors, w_t, has independent components and is obtained from the reduced-form errors, u_t, by a linear transformation, $u_t = Bw_t$. Under this assumption and the additional assumption that at most one component of w_t has a Gaussian distribution, the matrix B is unique up to column permutations and column signs. In some related approaches, specific distributions are assumed (see, e.g., Lanne et al. (2017)), while other ICA methods allow for more general non-Gaussian distribution families (see Hyvärinen et al. (2001), Gouriéroux et al. (2017), Hafner et al. (2025)).

If B is available, structural impulse responses can be computed from the VAR slope coefficients, $A_1, \ldots, A_p$, and B as follows:

$$\Theta_j = \Phi_j B, \quad j = 0, 1, \ldots,$$

where $\Phi_j = \sum_{i=1}^{j} \Phi_{j-i} A_i, j = 1, 2, \ldots,$ with $\Phi_0 = I_K$ and $A_i = 0$ for $i > p$ (e.g., Lütkepohl (2005, Section 2.1.2)). Thus, the shock transmission will be time-varying if only B varies over time even if the VAR slope coefficients are time-invariant.

2.2 ML Estimation of Structural Parameters

Lanne et al. (2017) assume that the structural errors, w_{kt}, have zero mean and finite variance σ_k^2, they are mutually stochastically independent and at most one of them has a Gaussian distribution. The density of w_{kt} is denoted by $\sigma_k^{-1} f_k(x/\sigma_k, \eta_k)$. In other words, the structural errors may have distributions from different distribution families, although they may of course be from the same distribution family and only differ in the specific parameter values or may even have identical parameters. The log likelihood function of the structural VAR model for a sample of size T can then be set up as

$$L_T(\theta) = \frac{1}{T} \sum_{t=1}^{T} l_t(\theta), \tag{3}$$

where

$$l_t = \sum_{k=1}^{K} \log f_k(\sigma_k^{-1} \iota_K' B(\beta)^{-1} u_t(\alpha); \eta_k) - \log|\det B(\beta)| - \sum_{k=1}^{K} \log \sigma_k$$

(see Eq. (7)/(8) of Lanne et al. (2017)). Here $\iota_K = (1, \dots, 1)'$ is a K-dimensional vector of ones, $\alpha = \text{vec}(A_1, \dots, A_p)$ and β contains the free parameters of B. Lanne et al. parameterize the B matrix so as to make the column sign and column order unique. Then by maximizing this log-likelihood, the ML estimator of the structural parameters B is obtained. Clearly, this approach starts out from independent structural errors and thereby ensures that a unique structural matrix B exists. In addition to this assumption, also specific distributions or at least distribution families for the structural errors w_{kt} have to be assumed. Under these assumptions, asymptotic properties can be derived (for details see Lanne et al. (2017)).

2.3 Pseudo ML Approach

Gouriéroux et al. (2017) consider a pseudo ML (PML) approach which gets away without assuming knowledge of the specific distribution families for the structural errors. They assume that the reduced-form errors are standardized such that they have zero mean and identity covariance matrix. Let $u_t^* = \Sigma_u^{-1/2} u_t \sim (0, I_K)$ denote the standardized residuals. Then the choice of a transformation matrix B such that $w_t = B^{-1} u_t$ has independent components, amounts to choosing an orthogonal matrix $Q \in \mathcal{O}(K)$, $\mathcal{O}(K)$ being the set of orthogonal $(K \times K)$ matrices, such that $w_t = Q' u_t^*$ has independent components and defining $B = \Sigma_u^{1/2} Q$.

Gouriéroux et al. (2017) propose to set up a pseudo log-likelihood function,

$$\log l_T(Q) = \sum_{t=1}^{T} \sum_{k=1}^{K} \log g_k(q_k' u_t), \tag{4}$$

where $g_k(\cdot)$ $(k = 1, \dots, K)$ is a non-Gaussian probability density function (p.d.f.) which may differ from the true p.d.f. of w_{kt} and q_k is the k^{th} column of Q (i.e., q_k' is the k^{th} row of $Q' = Q^{-1}$). Maximizing $\log l_T(Q)$ with respect to Q subject to $Q'Q = I_K$, gives the PML estimator. The latter restriction reduces the parameter space to dimension $K(K-1)/2$. Thus, Q can be parameterized by a $K(K-1)/2$ dimensional vector θ, $Q = Q(\theta)$, and $\log l_T(Q(\theta))$ just has to be maximized over θ.

Note that an orthogonal matrix Q can be written as a product of Givens matrices,

$$Q(\theta) = \left(\prod_{k=1}^{K-1} \prod_{j=k+1}^{K} \mathcal{G}_{k,j,}(\theta_i) \right)', \tag{5}$$

where

$$\mathcal{G}_{k,j}(\theta_i) = [g_{ij}]$$

is a $(K \times K)$ Givens matrix such that for $k \neq j$, $g_{kk} = g_{jj} = \cos\theta_i$, $g_{kj} = \sin\theta_i$, and $g_{jk} = -\sin\theta_i$. All other elements g_{mn} are 0 for $m \neq n$ and 1 for $m = n$. For example, for $K = 3$,

$$\mathcal{G}_{1,3}(\theta_i) = \begin{bmatrix} \cos\theta_i & 0 & \sin\theta_i \\ 0 & 1 & 0 \\ -\sin\theta_i & 0 & \cos\theta_i \end{bmatrix}.$$

Moreover, $\theta_i \in (0, 2\pi]$ for $i = 1, \dots, K(K-1)/2$, such that $\log l_T(Q(\theta))$ can in principle be optimized by grid search. This may be quite feasible for small K but could be a challenge for larger K. Note that for $K = 5$, θ has dimension 10 already.

Gouriéroux et al. (2017) show that the PML estimator is consistent and asymptotically normal under quite general conditions for the true distributions and the assumed p.d.f.s $g_k(\cdot)$. Of course, in practice the unknown standardized errors u_t have to be replaced by estimators. Gouriéroux et al. (2017) propose to estimate the reduced-form VAR process by least-squares (LS) and use the standardized LS estimates $\widetilde{\Sigma}_u^{-1/2}\hat{u}_t$ in place of u_t^*. Here $\widetilde{\Sigma}_u = T^{-1}\sum_{t=1}^{T}\hat{u}_t\hat{u}_t'$ is the usual error covariance estimator based on the LS residuals $\hat{u}_t$ and $\widetilde{\Sigma}_u^{1/2}$ is such that $\widetilde{\Sigma}_u^{1/2}\widetilde{\Sigma}_u^{1/2\prime} = \widetilde{\Sigma}_u$. For example, $\widetilde{\Sigma}_u^{1/2}$ may be a Cholesky component of $\widetilde{\Sigma}_u$.

2.4 Other Approaches and Discussion

In our simulations and in the empirical example we will use the aforementioned methods. However, there are a number of alternative approaches that have been used in the structural VAR literature to estimate the structural parameters and the shocks. For example, Hafner et al. (2025) propose to estimate the distributions of the structural shocks nonparametrically by kernel methods and then maximize the estimated likelihood function obtained in this way with respect to the structural parameters.

Herwartz (2018) proposes to choose the structural errors from the reduced-form errors by maximizing the p-value of an independence test for the structural errors. In practice, a minimum requirement is that the maximal p-value exceeds some threshold value such as 5% or 10%. Otherwise identification of the shocks through non-Gaussianity is not ensured. The problem is, of course, that a statistical test not rejecting the independence assumption may not mean that the components are independent but only that there is not enough evidence to reject independence.

All these approaches work under the assumption that a linear transformation exists that turns the reduced-form residuals into independent components. Since there exist distributions for which this assumption is not valid, it is not surprising that the ICA tool kit is substantially larger than what we have reviewed here. As we are not using such alternative methods in the following, we refer readers to Hyvärinen et al. (2001), Matteson and Tsay (2017) and Hafner et al. (2025) for discussions of further tools.

Another basic assumption underlying our analysis is, of course, the non-Gaussianity of the structural shocks. That assumption is often investigated by considering the non-Gaussianity of the reduced-form errors u_t. Clearly, all components of u_t will generally be non-Gaussian, even if only one component of w_t is non-Gaussian and all other components are Gaussian. Thus, the statistical underpinning of the non-Gaussianity assumption for the structural shocks may be weak and it should be understood that it is often an assumed rather than a confirmed property.

In the aforementioned estimation procedures yet another crucial assumption is the time-invariance of the distributions of the structural shocks across the sample. If that assumption is violated, the impact effects matrix B may also change during the sample period and then the shock transmission will change even if the reduced-form slope coefficients $A_1, \dots, A_p$ are time-invariant. As the data should be informative about a time-varying B matrix, we will consider statistical tests for this type of time-varying shock transmission in the next section.

3 Testing for Time-Varying Shock Transmission

Time-varying shock transmission is possible if the VAR slope coefficients or the impact effects of the shocks, B, vary. In this section we continue to assume that the VAR slope parameters are time-invariant. Thus, changes in the shock transmission must be due to changes in B. As the identification of the latter matrix requires suitable identifying information, additional information may be needed when the impact effects are changing. If sufficient information is available, testing for changing impact effects becomes possible. Lütkepohl and Schlaak (2022) combine identifying information from external instruments and heteroskedasticity to assess time-variation in the impact effects while Bruns and Lütkepohl (2024) consider restrictions on the long-run effects of the shocks in combination with heteroskedasticity to test for time-varying impact effects. In the following we will show how identifying information from non-Gaussianity can be used to explore time-variation in the impact effects that can even be due to other features of the distribution than heteroskedasticity.

3.1 General Changes in the Shock Distribution

Suppose we suspect a change in the distribution of the shocks and a corresponding change in the shock transmission at the end of time period T_1. In other words, the structural parameters are $\theta^{(1)}$ for $t \in \mathcal{T}_1 = \{1, \ldots, T_1\}$ and $\theta^{(2)}$ for $t \in \mathcal{T}_2 = \{T_1 + 1, \ldots, T\}$. Then we would like to test

$$\mathbb{H}_0 : \theta^{(1)} = \theta^{(2)} \quad \text{vs.} \quad \mathbb{H}_1 : \theta^{(1)} \neq \theta^{(2)}. \tag{6}$$

As the shocks are assumed to be instantaneously independent, it may be reasonable to assume that the u_t are not only serially uncorrelated but even serially independent. Since the structural parameters are parameters of the distribution of the residuals, estimates of $\theta^{(1)}$ and $\theta^{(2)}$ are asymptotically independent, assuming, e.g., that $\theta^{(m)}$ is estimated based on $\log l_T(Q(\theta))$ for the sample period $\mathcal{T}_m$. Suppose

$$\sqrt{T}(\hat{\theta}^{(m)} - \theta^{(m)}) \xrightarrow{d} \mathcal{N}(0, \tau_m^{-1} V_m), \quad m = 1, 2, \tag{7}$$

where $\tau_1 = T_1/T$ and $\tau_2 = (T - T_1)/T$ are the fractions of the sample associated with $\mathcal{T}_1$ and $\mathcal{T}_2$, respectively. These fractions are assumed to be independent of T. Under this assumption, the asymptotic distribution in (7) is obtained from results in Lanne et al. (2017) or Gouriéroux et al. (2017) if their conditions are satisfied.

Based on the asymptotic result in (7) we can test the pair of hypotheses (6) using the test statistic

$$W = T(\hat{\theta}^{(1)} - \hat{\theta}^{(2)})' \left(\tau_1^{-1}\hat{V}_1 + \tau_2^{-1}\hat{V}_2\right)^{-1} (\hat{\theta}^{(1)} - \hat{\theta}^{(2)}). \tag{8}$$

Under the null hypothesis, this statistic has an asymptotic $\chi^2(K(K-1)/2)$ distribution. Here $\hat{V}_m$ is a suitable estimator of the asymptotic covariance matrix V_m. For this test to be feasible, we need estimates of V_m, $m = 1, 2$. We will use a bootstrap method to estimate V_m from $\mathcal{T}_m$ (see Appendix B.1 for the details).

This test would even work if the shocks identified by non-Gaussianity have no economic meaning and are therefore perhaps not even proper structural shocks. The Wald test based on the test statistic in Equation (8) tests time-varying impact effects of the structural impulse responses or some transformation of them. If the test rejects, it rejects time-invariance of the impact effects of the shocks and, hence, of the shock transmission, even if Σ_u is time-invariant. These considerations show that there can be a change in the shock transmission even if the reduced-form VAR parameters including the residual covariance Σ_u are time-invariant.

3.2 Heteroskedastic Shocks

Heteroskedasticity is a specific change in the distribution of the shocks which is often visible even in the reduced-form residuals. If such a change in the distribution is diagnosed or suspected, it may be preferable to take time-variation in Σ_u explicitly into account. In that case, one may want to test time-varying impact effects directly by considering $B = \Sigma_u^{1/2} Q(\theta)$. As it is often of interest to test time-varying impulse responses of a specific structural shock, we focus on the case of testing for time-varying impulse responses of a single structural shock. To investigate the time-invariance of all the impulse responses, a separate test can be performed for each shock.

To simplify the exposition we assume that interest focusses on testing time-variation of the first column of the B matrices related to $\mathcal{T}_m$, $m = 1, 2$. In the following we denote these matrices by $B^{(1)}$ and $B^{(2)}$ and their first columns by $b^{(1)}$ and $b^{(2)}$, respectively.

Let $\Sigma_u^{(m)}$ be the covariance matrix of $u_t, t \in \mathcal{T}_m$. Then $\Sigma_u^{(m)} = B^{(m)} B^{(m)\prime}$, if we continue to standardize the variances of the shocks to be one ($\Sigma_w = I_K$) across the full sample period. Thus, the $B^{(m)}$ matrices reflect possible changes in the impact effects of the shocks and changes in their variability. To separate these two elements of variability, it is now preferable to standardize one element in each column of $B^{(m)}$ to one (e.g., the elements on the main diagonal) and allow w_t to have variances different from one such that $\Sigma_u^{(m)} = B^{(m)}\Sigma_w^{(m)}B^{(m)\prime}$, $m = 1, 2$, where the $\Sigma_w^{(m)}$ are diagonal matrices. Thereby it becomes more apparent that changes in the reduced-form covariances can be due to changes in the impact effects of the shocks, the shock variances or both. Note, however, that for two covariance matrices $\Sigma_u^{(1)}$ and $\Sigma_u^{(2)}$ we can always find a matrix B such that $\Sigma_u^{(1)} = BB'$ and $\Sigma_u^{(2)} = B\Lambda B'$, where Λ is a diagonal matrix. Thus, we can always find a decomposition of the $\Sigma_u^{(m)}$ matrices that ensures uncorrelated shocks $w_t = B^{-1}u_t$ in both volatility regimes that have time-invariant impact effects B. They may not be independent, however, if the shocks are non-Gaussian. It is, of course, possible that B provides even independent shocks. In that case the change in the reduced-form covariance matrix must be exclusively due to changes in the variances of the structural shocks and the shock transmission is time-invariant.

If one element of each column of $B^{(m)}$ is standardized to one, variance changes of the shocks w_t are reflected in changes in $\Sigma_w^{(m)}$. Our standardization of $B^{(m)}$ implies that one element of $b^{(m)}$ is also one. Without loss of generality, we assume that the first element is one and denote the vector of the second to last elements of the normalized vector by $\beta^{(m)}$.[2] Then the goal is to test the pair of hypotheses

$$\mathbb{H}_0 : \beta^{(1)} = \beta^{(2)} \quad \text{vs.} \quad \mathbb{H}_1 : \beta^{(1)} \neq \beta^{(2)}. \tag{9}$$

An obvious estimator of $\beta^{(m)}$ would be

$$\hat{\beta}^{(m)} = [0, I_{K-1}]\hat{b}^{(m)}/\hat{b}_1^{(m)},$$

where $\hat{b}^{(m)}$ is the first column of $\hat{\Sigma}_u^{(m)1/2}Q(\hat{\theta}^{(m)})$ and $\hat{b}_1^{(m)}$ is its first component. If heteroskedasticity is a possibility, $\Sigma_u^{(m)}$ should be estimated from $\hat{u}_t, t \in \mathcal{T}_m$, of course.

If the estimator $\hat{\theta}^{(m)}$ is asymptotically normal as in (7), the same holds for $\hat{b}^{(m)}$ estimated in this way by appealing to Slutsky's Theorem, because $Q(\theta)$ is a differentiable function of θ. Thus, a suitable statistic for testing the pair of hypotheses in (9) is

$$W_1 = T(\hat{\beta}^{(1)} - \hat{\beta}^{(2)})' \left(\tau_1^{-1}\hat{V}_{\beta 1} + \tau_2^{-1}\hat{V}_{\beta 2}\right)^{-1} (\hat{\beta}^{(1)} - \hat{\beta}^{(2)}), \tag{10}$$

which has an asymptotic $\chi^2(K-1)$ distribution. We will use again bootstrap methods to construct suitable estimators of the covariance matrices $V_{\beta m}$, $m = 1, 2$ (see Appendix B.1 for details).

[2] If the first element of $\beta^{(m)}$ is zero, we can standardize the first nonzero element of $\beta^{(m)}$ instead and reorder the variables such the normalized element is the first one. Of course, the same variable ordering has to be used in both regimes to ensure comparable impact effects.

A drawback of this test may be that it has to be ensured that the first column of each of the $B^{(m)}$ matrices ($m = 1, 2$) contains the impact effects of the same shock. Recall that identification and estimation of the impact effects using purely statistical criteria such as the distributional properties implies that the labelling of the shocks can usually only be done after the estimation because the columns will be ordered arbitrarily. Hence, the first column of $B^{(1)}$ may not represent the impact effects of the same shock as that of $B^{(2)}$. In practice, that issue may create a problem and it is difficult to give general advise how to deal with it because the labelling of the shocks depends on the specific empirical model. Clearly, the test based on (10) is meant to test time-variation in the impact effects of the same shock. In Section 5, we discuss an example and show how the problem can be circumvented in a specific empirical study.

The discussion of this section has been in terms of two transmission regimes. Generalizing it for more regimes is straightforward. A critical issue is, however, the choice of regimes. This has not been discussed so far but should ideally be linked to subject matter considerations. For example, for a monetary policy analysis it may be linked to special events concerning monetary policy such as policy changes announced by the central bank. Alternatively, one may also consider statistical procedures to investigate volatility changes during the sample. Such an approach may lead to pretesting issues that we do not consider in this study.

4 Small Sample Properties of the Tests

To get an impression of the small sample rejection frequencies of the two tests proposed in Equations (8) and (10) under the null and alternative hypotheses, we have performed a Monte Carlo simulation. Test statistic (8) is appropriate when testing for general time-varying transmission of homoskedastic shocks, while test statistic (10) may be preferable for heteroskedastic shocks or if one is interested in the transmission of a specific shock. Recall that in the heteroskedastic case, the change in the distribution may be exclusively due to changing variances of the structural shocks, while leaving the transmission of the shocks time-invariant if shocks of a fixed size are considered. On the other hand, focussing the test on the impact effects of individual shocks, the issue of comparing the same shocks across different volatility regimes may create additional problems in practice. Since, in practice, the variance properties of the shocks may not be known, the test based on (10) may still be a good choice if the reduced-form residuals are heteroskedastic.

We analyze the test properties for structural VAR models with two and three variables, setting the VAR order to zero. Parameter estimation is performed with pseudo maximum likelihood as in Equation (4), where the assumed distributions $g_k(\cdot)\ \forall k$ follow t-distributions with 4 degrees of freedom. The simulation is carried out under the general setting of $R = 1\,000$ replications. One of the generated structural shocks follows a normal distribution, the remaining structural shocks have t-distributions. The t-distributions have 4 degrees of freedom for homoskedastic DGPs and 4 and 3 degrees of freedom for heteroskedastic DGPs.[3] We consider sample sizes of $T = 200, 400, 1\,000, 2\,000$ and investigate a single break in the shock transmission matrix B occurring at the break date $T_1 = T/2$. Exact parameter values for the DGPs are presented in Tables 7 and 8 in Appendix A. The variances of the estimators, V_m in Equation (8) and $V_{\beta m}$ in Equation (10), are estimated through a bootstrap with $S = 100$ replications carried out in each simulation run r. More details are given in Appendix B.1.

For both types of DGPs we consider alternative scenarios for the volatility of the reduced-form residuals and the structural shocks. In the first scenario, called the *homoskedastic scenario*, the structural shocks w_t are always homoskedastic while the reduced-form residuals u_t are homoskedastic under the null hypothesis of time-invariance of the impact effects but heteroskedastic under $\mathbb{H}_1$, where the impact effects are time-varying (see Table 7). In a second scenario, called the *heteroskedastic scenario*, the variances of both, the structural as well as the reduced-form residuals, are distinct in the two volatility regimes considered (see Table 8). In that case, only the impact effects are time-invariant under $\mathbb{H}_0$. Thus, under the null hypothesis, $\Sigma_u^{(m)}$ is time-varying only due to the time variation in $\Sigma_w^{(m)}$, whereas under the alternative hypothesis, $\Sigma_u^{(m)}$ is time-varying due to both, time variation in $\Sigma_w^{(m)}$ and time variation in the matrix of impact effects.

The simulation results are reported in Tables 1 - 4. The overall impression is that the asymptotic critical values from the proper χ^2-distributions are useful in small samples. The small sample rejection frequencies line up quite well with the nominal significance levels, α, in particular for the larger sample sizes of $T = 1\,000$ and $2\,000$, while the power is quite impressive. For the larger sample sizes, it is in fact one or close to one in all scenarios considered. Thus, the test has a good chance to find a structural break in larger samples. For the smaller sample sizes of $T = 200$ and 400, the tests tend to reject slightly too often under $\mathbb{H}_0$, which may contribute to the large rejection frequencies for the smaller sample sizes when $\mathbb{H}_0$ is false.

[3] Note that the variance of a t-distribution with d degrees of freedom is given by $\frac{d}{d-2}$.

Table 1: Size and Power of the Wald Test for Time-Varying Shock Transmission Based on Test Statistic (8) for a Bivariate DGP with *Homoskedastic* w_t

	Size ($\theta^{(1)} = \theta^{(2)}$)			Power ($\theta^{(1)} \neq \theta^{(2)}$)		
T	$\alpha = 0.1$	0.05	0.01	$\alpha = 0.1$	0.05	0.01
200	0.202	0.131	0.050	0.767	0.646	0.445
400	0.161	0.108	0.040	0.842	0.751	0.555
1000	0.089	0.062	0.023	0.964	0.939	0.832
2000	0.099	0.046	0.008	0.997	0.997	0.965

Notes: No. of replications $R = 1\,000$, variances of the estimator are estimated by a bootstrap with $S = 100$ replications in each simulation run. The nominal test size, α, is based on an asymptotic $\chi^2(1)$-distribution.

Table 2: Size and Power of the Wald Test for Time-Varying Shock Transmission Based on Test Statistic (8) for a Three-Dimensional DGP with *Homoskedastic* w_t

	Size ($\theta^{(1)} = \theta^{(2)}$)			Power ($\theta^{(1)} \neq \theta^{(2)}$)		
T	$\alpha = 0.1$	0.05	0.01	$\alpha = 0.1$	0.05	0.01
200	0.175	0.103	0.028	0.982	0.960	0.883
400	0.141	0.076	0.016	0.999	0.995	0.971
1000	0.126	0.063	0.015	1.000	1.000	0.996
2000	0.129	0.063	0.013	1.000	1.000	1.000

Notes: No. of replications $R = 1\,000$, variances of the estimator are estimated by a bootstrap with $S = 100$ replications in each simulation run. The nominal test size, α, is based on an asymptotic $\chi^2(3)$-distribution.

5 Empirical Example

As an empirical illustration, we apply the tests to the structural VAR model estimated in Kilian (2009) over the sample 1973M2 until 2007M12. Kilian (2009) analyzes the oil market in a recursively identified (Cholesky) structural VAR. The system includes three variables – percent change in global crude oil production ($y_{1t}^{\Delta\text{prod}}$), an index of real economic activity (y_{2t}^{rea}) and the real price of oil (y_{3t}^{rpo}) – and three structural shocks: an oil supply shock ($w_{1t}^{\text{oil supply}}$), an aggregate demand shock ($w_{2t}^{\text{aggregate demand}}$) and an oil specific-demand shock ($w_{3t}^{\text{oil-specific demand}}$). Kilian (2009) uses a VAR order of $p = 24$. For easy comparison of our results, we also use $p = 24$.

Using our notation in Equation (1), the system of Kilian (2009) can be

Table 3: Size and Power of the Wald Test for Time-Varying Shock Transmission Based on Test Statistic (10) for a Bivariate DGP with *Heteroskedastic* w_t

	Size ($\beta^{(1)} = \beta^{(2)}$)			Power ($\beta^{(1)} \neq \beta^{(2)}$)		
T	$\alpha = 0.1$	0.05	0.01	$\alpha = 0.1$	0.05	0.01
200	0.143	0.088	0.027	0.893	0.838	0.712
400	0.127	0.076	0.023	0.951	0.927	0.839
1000	0.122	0.069	0.023	0.994	0.992	0.978
2000	0.109	0.056	0.015	1.000	1.000	1.000

Notes: No. of replications $R = 1\,000$, variances of the estimator are estimated by a bootstrap with $S = 100$ replications in each simulation run. The nominal test size, α, is based on an asymptotic $\chi^2(1)$-distribution. β refers to the first column of B. The log likelihood accounts for heteroskedasticity.

Table 4: Size and Power of the Wald Test for Time-Varying Shock Transmission Based on Test Statistic (10) for a Three-Dimensional DGP with *Heteroskedastic* w_t

	Size ($\beta^{(1)} = \beta^{(2)}$)			Power ($\beta^{(1)} \neq \beta^{(2)}$)		
T	$\alpha = 0.1$	0.05	0.01	$\alpha = 0.1$	0.05	0.01
200	0.191	0.123	0.047	0.943	0.910	0.834
400	0.159	0.101	0.035	0.982	0.970	0.929
1000	0.109	0.077	0.020	1.000	1.000	0.992
2000	0.110	0.071	0.018	1.000	1.000	1.000

Notes: No. of replications $R = 1\,000$, variances of the estimator are estimated by a bootstrap with $S = 100$ replications in each simulation run. The nominal test size, α, is based on an asymptotic $\chi^2(2)$-distribution. β refers to the first column of B. The log likelihood accounts for heteroskedasticity.

represented as

$$u_t \equiv \begin{pmatrix} u_{1t}^{\Delta\text{prod}} \\ u_{2t}^{\text{rea}} \\ u_{3t}^{\text{rpo}} \end{pmatrix} = \begin{pmatrix} b_{11} & 0 & 0 \\ b_{21} & b_{22} & 0 \\ b_{31} & b_{32} & b_{33} \end{pmatrix} \begin{pmatrix} w_{1t}^{\text{oil supply}} \\ w_{2t}^{\text{aggregate demand}} \\ w_{3t}^{\text{oil specific-demand}} \end{pmatrix}. \quad (11)$$

The estimated matrix of impact effects of the shocks is

$$\hat{B}_{rec} = \begin{pmatrix} -18.74 & 0 & 0 \\ -0.07 & 4.06 & 0 \\ 0.44 & -0.47 & 5.94 \end{pmatrix},$$

where the shock variances are normalized to one ($\Sigma_w = I_3$). For comparison, in a non-Gaussian structural VAR, we estimate the following impact effects,

$$\hat{B}_{nG} = \begin{pmatrix} -18.60 & 2.24 & -0.25 \\ 0.41 & 4.03 & 0.31 \\ 0.36 & -0.03 & 5.96 \end{pmatrix}.$$

In particular, the diagonal elements are close to those from Kilian's study, suggesting that the ordering of the shocks can be assumed to be the same as in Equation (11).

The corresponding reduced-form residuals of Kilian's model are shown in Figure 1. A heteroskedastic pattern is clearly visible. We interpret that pattern as indication that the distribution of the shocks has changed during the sample period. This could be due to hetereskedasticity of the structural shocks only while the transmission matrix B remains unchanged. However, the heteroskedasticity could also indicate that the impact effects matrix B is time-varying. We use our tests proposed in Section 3 to assess this possibility.

Applying the tests for structural breaks requires to choose a break date. Barsky and Kilian (2004) discuss turbulences in the oil market associated with Iraq's invasion of Kuwait in 1990. Bruns and Lütkepohl (2023) test for variance changes, providing evidence for a volatility regime change in 1990M9. In line with their finding, we split the sample in 1990M9 to illustrate our tests, although further volatility changes may be present.

Assuming a potential change in the impact effects of the shocks in 1990M9 we get the following subsample estimates of the $B^{(m)}$ matrices, using a recursive structural VAR,

$$\hat{B}_{rec}^{(1)} = \begin{pmatrix} -23.90 & 0 & 0 \\ 0.24 & 4.62 & 0 \\ 0.61 & 0.83 & 4.75 \end{pmatrix} \text{ and } \hat{B}_{rec}^{(2)} = \begin{pmatrix} -8.84 & 0 & 0 \\ -0.29 & 3.34 & 0 \\ -0.35 & 1.01 & 6.22 \end{pmatrix}$$

and, using a non-Gaussian SVAR,

$$\hat{B}_{nG}^{(1)} = \begin{pmatrix} -13.91 & 12.95 & 14.49 \\ -0.47 & 2.98 & -3.51 \\ 4.06 & 2.62 & 0.54 \end{pmatrix} \text{ and } \hat{B}_{nG}^{(2)} = \begin{pmatrix} -7.76 & -0.57 & 4.19 \\ -1.36 & 2.64 & -1.55 \\ 1.50 & 4.52 & 4.14 \end{pmatrix}.$$

In the non-Gaussian SVAR, the shocks have to be labelled properly. We do so by aligning the order and the signs of the columns of $\hat{B}_{nG}^{(1)}$ and $\hat{B}_{nG}^{(2)}$ with the signs proposed by Kilian and Murphy (2012). These authors use sign restrictions to identify the shocks and impose the sign pattern

$$\begin{pmatrix} - & + & + \\ - & + & - \\ + & + & + \end{pmatrix} \tag{12}$$

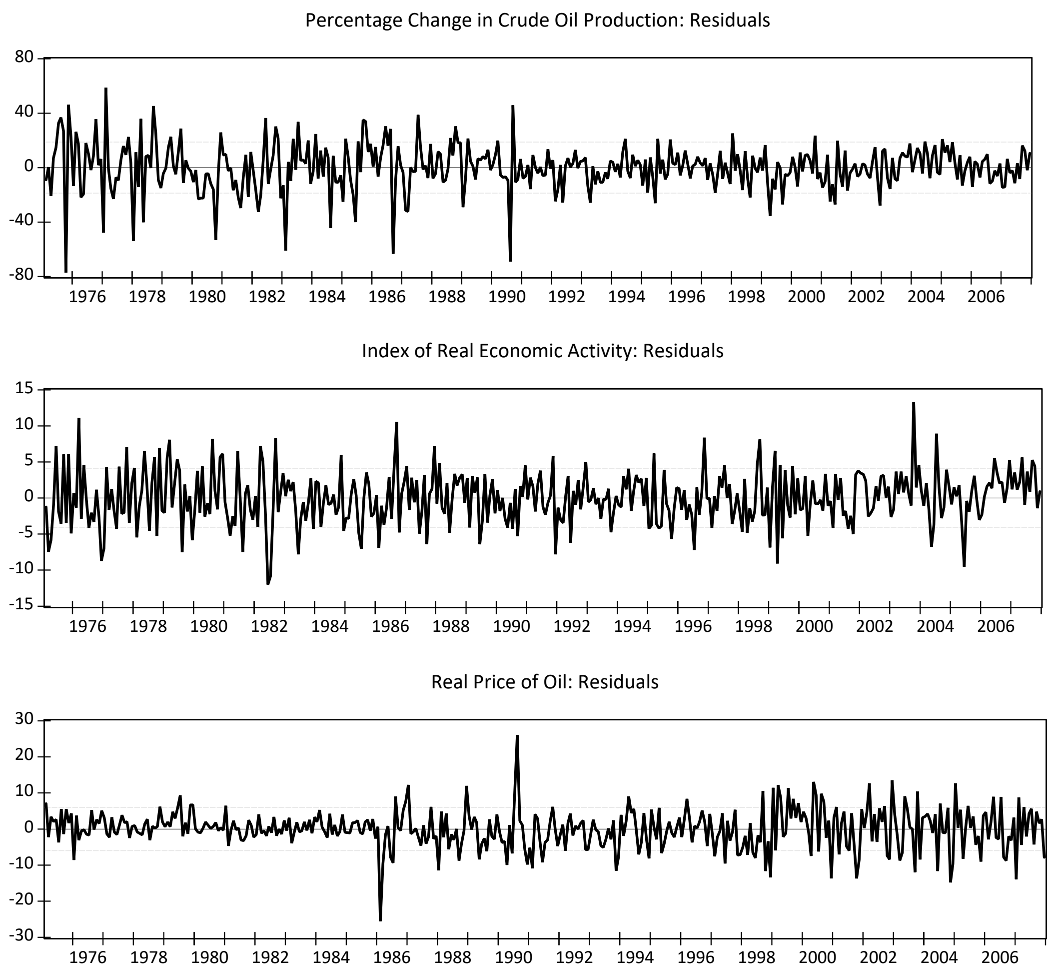

Figure 1: Reduced-form LS residuals of the VAR(24) model of the crude oil market.

Table 5: Tests for Time-Varying Impact Effects Based on Test Statistic (8)

$\mathbb{H}_0$	Change Point	Wald Statistic	p-value
$\theta^{(1)} = \theta^{(2)}$	1990M7	47.441	0.000
	1990M8	16.410	0.001
	1990M9	16.620	0.001
	1990M10	17.522	0.001
	1990M11	20.200	0.000

Note: The V_m are estimated using a bootstrap as described in Appendix B.1 with $S = 2\,000$ replications. The p-values are based on a $\chi^2(3)$-distribution.

Table 6: Tests for Time-Varying Impact Effects Based on Test Statistic (10)

	oil supply		aggregate demand		oil-specific demand	
Change in	Wald Stat.	p-value	Wald Stat.	p-value	Wald Stat.	p-value
1990M7	2.266	0.322	2.562×10^3	0.000	143.994	0.000
1990M8	0.525	0.769	1.289×10^3	0.000	20.753	0.000
1990M9	0.511	0.774	3.334×10^3	0.000	7.284	0.026
1990M10	0.506	0.777	0.718×10^3	0.000	13.254	0.001
1990M11	0.571	0.756	3.678×10^3	0.000	12.743	0.002

Note: The $V_{\beta m}$ are estimated using a bootstrap as described in Appendix B.1 with $S = 2\,000$ replications. The p-values are based on a $\chi^2(2)$-distribution.

on the impact effects matrix. We have used that pattern as a basis for ordering the columns of $\hat{B}_{nG}^{(1)}$ and $\hat{B}_{nG}^{(2)}$. As the estimates in the columns of these two matrices look quite different, testing formally for time-varying impact effects may be worthwhile.

First, we apply the test based on (8) that standardizes the structural shocks to be homoskedastic and all heteroskedasticity is due to a change in the impact effects. Note that if the true data generating process features homoskedastic structural shocks but exhibits changes in the impacts effects matrix B, this results in heteroskedastic reduced-form residuals. The test results are presented in Table 5 for a set of possible change points around 1990M9. The test for change point 1990M9 clearly rejects time-invariance of the impact effects of the shocks as its p-value is less than 1%. Interestingly, the p-values for neighbouring months are also very small. Overall the indication for a time-varying impact effects matrix is quite robust to a potential slight misspecification of the actual change point.

Given that the change in the distribution may well be due to a change in the variances of the shocks and because we are interested to know whether the impact effects of all shocks are time-varying, we have also applied the test based on (10). Note that if the true data generating process features heteroskedastic structural shocks but a time-invariant B matrix, an estimated B matrix over two subsamples can seemingly appear to feature time-variation while in fact the change in B only picks up the true, but ignored change in the variances of the structural shocks. Applying the test based on (10), we obtain the results presented in Table 6. The p-values of the tests for time-varying impact effects of the oil supply shock are larger than 30% for all change dates considered. Thus, our test does not provide statistical evidence for time-varying impact effects of the oil supply shock. In contrast, the evidence for time-varying impact effects of the aggregate demand and the oil-specific demand shocks is quite strong. Almost all p-values for all assumed change dates are well below 1%. Thus, we can conclude that estimating the non-Gaussian model over each subsample separately is clearly supported by the data because at least some of the impact effects have changed during the sample period. It is again interesting to note that the test results are robust to slight changes in the assumed month of volatility change. The p-values obtained for July to November 1990 are very similar.

Overall the results of both types of tests in Tables 5 and 6 are well in line. They both present evidence of a change in the impact effects of the shocks in the second half of 1990. The second test permits a more detailed analysis of the impact effects of the individual shocks. It indicates that the change in impact effects may be primarily due to changes in the impact effects of the aggregate demand and oil-specific demand shocks, while there is no evidence of changes in the responses to an oil supply shock.

Figure 2 shows the estimated impulse response functions corresponding to shocks of size one-standard deviation when we allow for the impact effects $B^{(1)}$ and $B^{(2)}$ to be distinct. In the figure we also depict one-standard error bands for the impulse responses.[4] For comparison, we also show point estimates of the impulse responses when we identify triangular $\hat{B}^{(1)}_{rec}$ and $\hat{B}^{(2)}_{rec}$ from a recursive SVAR. In general, the impulse responses change, compared to the case of time-invariant impact effects in Kilian (2009).[5] However, the confidence intervals of the impact effects of the oil supply shock overlap, reflecting the insignificant test results in Table 6. In contrast, the impact effects of aggregate demand and oil-market specific demand shocks on oil production in the two subsamples appear to be quite different. Although the pre- and

[4] In line with Kilian (2009), the confidence intervals are calculated using a recursive design wild bootstrap with 2 000 replications as in Gonçalves and Kilian (2004), see also Appendix B.2.

[5] Interestingly, the two impulse responses from the recursively identified SVAR are also partly quite different.

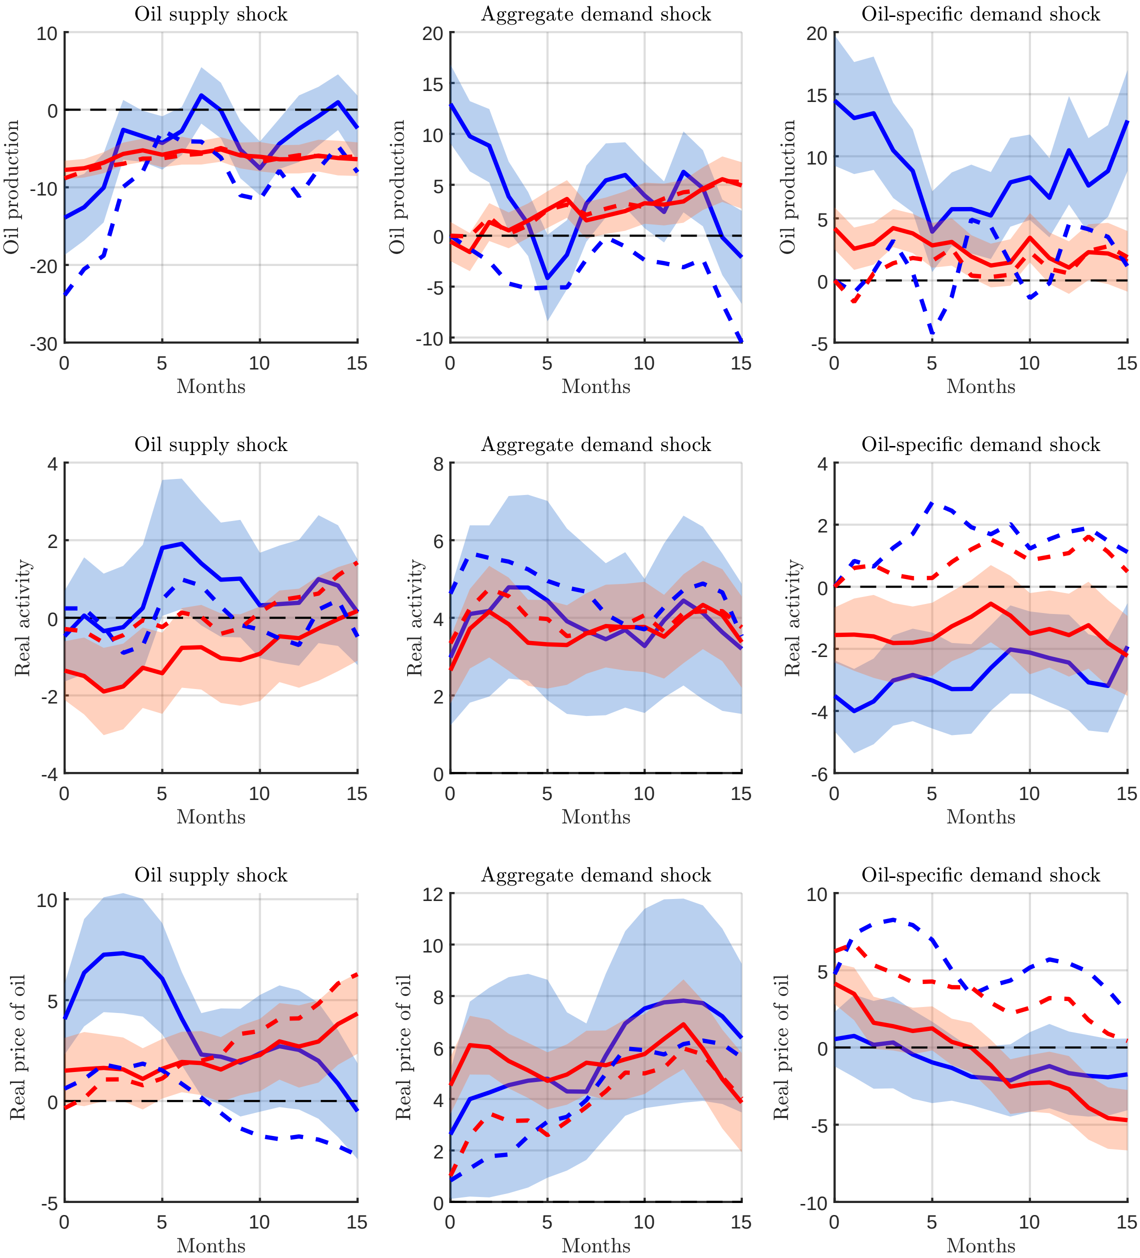

Figure 2: Comparison of impulse response functions. The solid lines show point estimates from non-Gaussian SVARs, the shaded areas represent one-standard error bands. Blue: pre-1990M9; red: post-1990M9. The dashed lines show point estimates from recursive SVARs.

post-1990M9 impact effects of the aggregate demand shocks on real activity are roughly the same and, hence, the size of the shocks is very similar, the effect on oil production is very different. The pre-1990M9 effect is much larger on impact than the post-1990M9 effect, where oil production increases only gradually after the shock has hit. The one-standard error oil specific demand shock is smaller pre-1990M9 but has a much larger effect on oil production (see the panels in the last column of Figure 2). Thus, it clearly makes a difference whether a possible change in the shock transmission is allowed for or not.

Note, however, that our interpretation of the shock transmission depends on our labelling of the shocks. The labelling of the last two shocks can, in fact, be questioned. Recall that the signs of the aggregate demand shock in the pre-1990M9 period are not fully in line with those assumed by Kilian and Murphy (2012) (see the sign matrix in (12)). Moreover, the shock labelled as 'oil-specific demand shock' has impact effects around zero on the real price of oil, as seen in Figure 2. Clearly that makes it difficult to think of the shock as a shock to the oil market. It is, of course, possible that a purely statistical identification, as in our identification strategy using non-Gaussianity, shocks are obtained without economic interpretation. Thus, our finding of a possible change in the shock transmission around 1990M9 may also reveal that the shocks obtained by the statistical identification procedure may not correspond to the economic shocks of interest.

6 Conclusions

Identifying causal relations is the centerpiece of structural VAR analysis. Identification through non-Gaussianity represents a statistical, data-driven approach. Another prominent statistical technique is identification through heteroskedasticity. Both approaches are often seen in isolation. We argue that it is important to see both techniques together as what they are: ways of exploiting the statistical properties of the data for structural identification. We also argue that all properties of the data should be taken into account. Heteroskedasticity is a change in the variances and, hence, in the distribution. The true data generating process then might in fact only feature changing variances but a time-invariant shock transmission matrix. Note that this is a crucial assumption in identification through heteroskedasticity. However, the true data generating process might as well exhibit a change in the shock transmission.

When identification is achieved through non-Gaussianity, time-varying shock transmission can and should be tested. In this paper, we propose two simple asymptotic χ^2-tests to do so. One test version is designed for an overall change in the shock transmission and the other one allows to test time-varying transmission for a specific shock and is particularly useful for heteroskedastic shocks. We have shown by simulation that the tests exhibit good small sample properties and we have applied them to an existing empirical study where time-invariant shock transmission has been assumed although the reduced-form residuals are heteroskedastic. We reject the null hypothesis of a time-invariant transmission matrix. Estimating a model that allows for time-varying impact effects of the shocks, our empirical example illustrates that the impulse responses along with their economic interpretation may change over time.

In this study, we have assumed that the possible change point of the distribution is known to the analyst, for example, from subject matter knowledge. In practice, the change point is likely to be uncertain. Although in our empirical example the test results were robust to changes in the assumed change point, it would be preferable to have statistical tools to aid in the determination of possible change points. If the change in the distribution is due to changes in the second moments, then there are tests for heteroskedasticity that may be suitable in the present context. For future research, it may be of interest, however, to develop tools that can be used more generally.

A Detailed Monte Carlo Settings

Table 7: Parameter Settings for a DGP with *homoskedastic* w_t

		bivariate DGP	three-dimensional DGP
$\mathbb{H}_0$	trans-mission matrix	$B^{(1)} = B^{(2)} = \begin{pmatrix} 1.0 & -1.7 \\ 2.0 & 1.0 \end{pmatrix}$ $\theta^{(1)} = \theta^{(2)} = -0.30$	$B^{(1)} = B^{(2)} = \begin{pmatrix} 1.0 & -1.7 & -0.9 \\ 2.0 & 1.0 & -3.5 \\ 1.5 & -1.3 & 1.0 \end{pmatrix}$ $\theta_1^{(1)} = \theta_1^{(2)} = 2.68$, $\theta_2^{(1)} = \theta_2^{(2)} = 0.33$, $\theta_3^{(1)} = \theta_3^{(2)} = 2.28$
	covariance matrix of structural shocks w_t	$\Sigma_w^{(1)} = \Sigma_w^{(2)} = \begin{pmatrix} 64 & 0 \\ 0 & 2 \end{pmatrix}$	$\Sigma_w^{(1)} = \Sigma_w^{(2)} \begin{pmatrix} 64 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix}$
	covariance matrix of reduced-form residuals u_t	$\Sigma_u^{(1)} = \Sigma_u^{(2)} = \begin{pmatrix} 70.0 & 124.5 \\ 124.5 & 258.0 \end{pmatrix}$	$\Sigma_u^{(1)} = \Sigma_u^{(2)} = \begin{pmatrix} 71.5 & 130.5 & 98.8 \\ 130.5 & 282.0 & 182.5 \\ 98.8 & 182.5 & 149.4 \end{pmatrix}$
$\mathbb{H}_1$	trans-mission matrix	$B^{(1)} = \begin{pmatrix} 1.0 & -1.7 \\ 2.0 & 1.0 \end{pmatrix}$ $B^{(2)} = \begin{pmatrix} 1.0 & -5.2 \\ 3.0 & 1.0 \end{pmatrix}$ $\theta^{(1)} = -0.30$ $\theta^{(2)} = -0.74$	$B^{(1)} = \begin{pmatrix} 1.0 & -1.7 & -0.9 \\ 2.0 & 1.0 & -3.5 \\ 1.5 & -1.3 & 1.0 \end{pmatrix}$ $B^{(2)} = \begin{pmatrix} 1.0 & -5.2 & -0.9 \\ 3.0 & 1.0 & -3.5 \\ 1.5 & -1.3 & 1.0 \end{pmatrix}$ $\theta_1^{(1)} = 2.68$, $\theta_2^{(1)} = 0.33$, $\theta_3^{(1)} = 2.28$ $\theta_1^{(2)} = 2.98$, $\theta_2^{(2)} = 0.75$, $\theta_3^{(2)} = 2.85$
	covariance matrix of structural shocks w_t	$\Sigma_w^{(1)} = \Sigma_w^{(2)} = \begin{pmatrix} 64 & 0 \\ 0 & 2 \end{pmatrix}$	$\Sigma_w^{(1)} = \Sigma_w^{(2)} = \begin{pmatrix} 64 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix}$
	covariance matrices of reduced-form residuals u_t	$\Sigma_u^{(1)} = \begin{pmatrix} 70.0 & 124.5 \\ 124.5 & 258.0 \end{pmatrix}$ $\Sigma_u^{(2)} = \begin{pmatrix} 118.0 & 181.6 \\ 181.6 & 578.0 \end{pmatrix}$	$\Sigma_u^{(1)} = \begin{pmatrix} 71.5 & 130.5 & 98.8 \\ 130.5 & 282.0 & 182.5 \\ 98.8 & 182.5 & 149.4 \end{pmatrix}$ $\Sigma_u^{(2)} = \begin{pmatrix} 119.5 & 187.6 & 107.8 \\ 187.6 & 602.0 & 278.5 \\ 107.8 & 278.5 & 149.4 \end{pmatrix}$

Note: The parameters for the bivariate case in the first column (upper and lower part) correspond to the simulation results reported in Table 1. The parameters for the three-dimensional case in the second column (upper and lower part) correspond to the simulation results reported in Table 2.

Table 8: Parameter Settings for a DGP with *heteroskedastic* w_t

		bivariate DGP	three-dimensional DGP
$\mathbb{H}_0$	trans-mission matrix	$B^{(1)} = B^{(2)} = \begin{pmatrix} 1.0 & -1.7 \\ 2.0 & 1.0 \end{pmatrix}$ $\theta^{(1)} = -0.64$ $\theta^{(2)} = -0.30$	$B^{(1)} = B^{(2)} = \begin{pmatrix} 1.0 & -1.7 & -0.9 \\ 2.0 & 1.0 & -3.5 \\ 1.5 & -1.3 & 1.0 \end{pmatrix}$ $\theta_1^{(1)} = 2.68,\ \theta_2^{(1)} = 0.70,\ \theta_3^{(1)} = 2.18$ $\theta_1^{(2)} = 2.68,\ \theta_2^{(2)} = 0.33,\ \theta_3^{(2)} = 2.28$
	covariance matrices of structural shocks w_t	$\Sigma_w^{(1)} = \begin{pmatrix} 16 & 0 \\ 0 & 3 \end{pmatrix}$ $\Sigma_w^{(2)} = \begin{pmatrix} 64 & 0 \\ 0 & 2 \end{pmatrix}$	$\Sigma_w^{(1)} = \begin{pmatrix} 16 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{pmatrix}$ $\Sigma_w^{(2)} = \begin{pmatrix} 64 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix}$
	covariance matrices of reduced-form residuals u_t	$\Sigma_u^{(1)} = \begin{pmatrix} 25.0 & 26.8 \\ 26.8 & 67.0 \end{pmatrix}$ $\Sigma_u^{(2)} = \begin{pmatrix} 70.0 & 124.5 \\ 124.5 & 258.0 \end{pmatrix}$	$\Sigma_u^{(1)} = \begin{pmatrix} 27.3 & 35.8 & 28.2 \\ 35.8 & 103.0 & 33.7 \\ 28.2 & 33.7 & 44.1 \end{pmatrix}$ $\Sigma_u^{(2)} = \begin{pmatrix} 71.5 & 130.5 & 98.8 \\ 130.5 & 282.0 & 182.5 \\ 98.8 & 182.5 & 149.4 \end{pmatrix}$
$\mathbb{H}_1$	trans-mission matrices	$B^{(1)} = \begin{pmatrix} 1.0 & -1.7 \\ 2.0 & 1.0 \end{pmatrix}$ $B^{(2)} = \begin{pmatrix} 1.0 & -2.6 \\ 3.0 & 1.0 \end{pmatrix}$ $\theta^{(1)} = -0.64$ $\theta^{(2)} = -0.43$	$B^{(1)} = \begin{pmatrix} 1.0 & -1.7 & -0.9 \\ 2.0 & 1.0 & -3.5 \\ 1.5 & -1.3 & 1.0 \end{pmatrix}$ $B^{(2)} = \begin{pmatrix} 1.0 & -2.6 & -0.9 \\ 3.0 & 1.0 & -3.5 \\ 1.5 & -1.3 & 1.0 \end{pmatrix}$ $\theta_1^{(1)} = 2.68,\ \theta_2^{(1)} = 0.70,\ \theta_3^{(1)} = 2.18$ $\theta_1^{(2)} = 2.68,\ \theta_2^{(2)} = 0.33,\ \theta_3^{(2)} = 2.51$
	covariance matrices of structural shocks w_t	$\Sigma_w^{(1)} = \begin{pmatrix} 16 & 0 \\ 0 & 3 \end{pmatrix}$ $\Sigma_w^{(2)} = \begin{pmatrix} 64 & 0 \\ 0 & 2 \end{pmatrix}$	$\Sigma_w^{(1)} = \begin{pmatrix} 16 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{pmatrix}$ $\Sigma_w^{(2)} = \begin{pmatrix} 64 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix}$
	covariance matrices of reduced-form residuals u_t	$\Sigma_u^{(1)} = \begin{pmatrix} 25.0 & 26.8 \\ 26.8 & 67.0 \end{pmatrix}$ $\Sigma_u^{(2)} = \begin{pmatrix} 77.5 & 186.8 \\ 186.8 & 578.0 \end{pmatrix}$	$\Sigma_u^{(1)} = \begin{pmatrix} 27.3 & 35.8 & 28.2 \\ 35.8 & 103.0 & 33.7 \\ 28.2 & 33.7 & 44.1 \end{pmatrix}$ $\Sigma_u^{(2)} = \begin{pmatrix} 71.5 & 194.5 & 98.8 \\ 194.5 & 602.0 & 278.5 \\ 98.8 & 278.5 & 149.4 \end{pmatrix}$

Note: The parameters for the bivariate case in the first column (upper and lower part) correspond to the simulation results reported in Table 3. The parameters for the three-dimensional case in the second column (upper and lower part) correspond to the simulation results reported in Table 4.

B Bootstraps

B.1 Bootstraps for Covariance Matrices

The estimates of the covariance matrices V_m and $V_{\beta m}$ used in the Wald statistics (8) and (10) of the tests for time-varying impact effects are estimated by the following bootstrap procedures.

For V_m, a large number, S, of estimates of the parameters $\theta^{(m)}$ are generated with a standard residual-based recursive design bootstrap based on random draws from the reduced-form LS residuals $\hat{u}_t$. Then we use the QML estimation procedure to get bootstrap estimates $\hat{\theta}_s^{(m)}$, $s = 1, \dots, S$. In this procedure it has to be ensured, however, that the ordering and signs of the columns of $Q(\theta)$ are the same in each bootstrap replication. Therefore, in the simulation exercise, the ordering and signs of the columns of $Q(\hat{\theta}_s^{(m)})$ follow the true $Q(\theta^{(m)})$-matrix of the data generating process. In the empirical example, the point estimate $Q(\hat{\theta}^{(m)})$ is taken as reference instead. Technically this is implemented by choosing the permutation of $Q(\hat{\theta}_s^{(m)})$ that has the highest correlation with the true $Q(\theta^{(m)})$ matrix in the Monte Carlo simulations and the estimate $Q(\hat{\theta}^{(m)})$ in the empirical example. Note that we also have to maximize the correlation over all combinations of column signs (see also below). Based on the $\hat{\theta}_s^{(m)}$ obtained in this way, the covariance matrix V_m is estimated in the usual way as

$$\hat{V}_m = \frac{1}{S-1} \sum_{s=1}^{S} \left(\hat{\theta}_s^{(m)} - \overline{\hat{\theta}^{(m)}} \right) \left(\hat{\theta}_s^{(m)} - \overline{\hat{\theta}^{(m)}} \right)', \text{ where } \overline{\hat{\theta}^{(m)}} = \frac{1}{S} \sum_{s=1}^{S} \hat{\theta}_s^{(m)}.$$

For $V_{\beta m}$, we proceed with the bootstrap as before and compute S bootstrap estimates $\hat{B}_s^{(m)}$, $s = 1, \dots, S$, of $B^{(m)}$. In this case, it has to be ensured that the ordering and signs of the shocks (i.e., the columns of $\hat{B}_s^{(m)}$) are the same in each bootstrap replication. We use the ordering and signs of the columns of $\hat{B}_s^{(m)}$ such that they follow the true $B^{(m)}$-matrix of the data generating process. In the empirical example, the point estimate $\hat{B}^{(m)}$ is taken as reference instead. Technically this is implemented as described for the $Q(\theta)$ (see above). Note that also all combinations of column signs have to be taken into account. Thus, for example for a three-dimensional system, we have to consider all permutations of $B = [b_1, b_2, b_3]$, $-B$, $[-b_1, b_2, b_3]$, $[b_1, -b_2, b_3]$, $[b_1, b_2, -b_3]$, $[-b_1, -b_2, b_3]$, $[-b_1, b_2, -b_3]$, and $[b_1, -b_2, -b_3]$. Here the b_i denote the columns of B. The correlation coefficient is calculated over all elements in the vectorized $\hat{B}_s^{(m)}$ and the vectorized $B^{(m)}$ matrix. Then we compute the first columns $\hat{\beta}_s^{(m)}$, $s = 1, \dots, S$, of the resulting standardized bootstrap estimates. Finally, the $V_{\beta m}$ matrix is estimated as

$$\hat{V}_{\beta m} = \frac{1}{S-1} \sum_{s=1}^{S} \left(\hat{\beta}_s^{(m)} - \overline{\hat{\beta}^{(m)}} \right) \left(\hat{\beta}_s^{(m)} - \overline{\hat{\beta}^{(m)}} \right)', \text{ where } \overline{\hat{\beta}^{(m)}} = \frac{1}{S} \sum_{s=1}^{S} \hat{\beta}_s^{(m)}.$$

In the Monte Carlo simulations of Section 4, bootstraps of this kind are performed in each replication.

B.2 Bootstrap for Impulse Responses

The bootstrap for the confidence intervals round the impulse responses is a wild bootstrap on the LS residuals of the reduced-form model as described in Kilian and Lütkepohl (2017, Section 12.2.3). In other words, the LS residuals $\hat{u}_t$ are replaced by $\eta_t \hat{u}_t$, where η_t is an independent random variable η_t with a standard normal distribution. In each bootstrap replication, the B matrix is estimated based on the pre- or post-1990M9 data and then its columns are permuted as described in Section B.1 such that the highest correlation with the point estimates $\hat{B}^{(1)}$ and $\hat{B}^{(2)}$ is obtained. In other words, in each bootstrap replication, the estimate that has the highest correlation with the original estimate is used for computing the bootstrap impulse responses from which the pointwise confidence intervals are constructed in the usual way using the percentiles of the bootstrap distributions.

References

Barsky, R. B. and Kilian, L. (2004). Oil and the macroeconomy since the 1970s, *Journal of Economic Perspectives* **18**: 115–134.

Bruns, M. and Lütkepohl, H. (2023). Have the effects of shocks to oil price expectations changed?: Evidence from heteroskedastic proxy vector autoregressions., *Economics Letters* **233**: 111416.

Bruns, M. and Lütkepohl, H. (2024). Heteroskedastic structural vector autoregressions identified via long-run restrictions, *Discussion Paper 2103*, DIW Berlin.

Comon, P. (1994). Independent component analysis, A new concept?, *Signal Processing* **36**: 287–314.

Comon, P., Jutten, C. and Herault, J. (1991). Blind separation of sources, Part ii: Problems statement, *Signal Processing* **24**: 11–20.

Gonçalves, S. and Kilian, L. (2004). Bootstrapping autoregressions with conditional heteroskedasticity of unknown form., *Journal of Econometrics* **123**: 89–120.

Gourieroux, C. and Monfort, A. (2014). Revisiting identification in structural VAR models, *Technical report*, CREST and University of Toronto.

Gouriéroux, C., Monfort, A. and Renne, J.-P. (2017). Statistical inference for independent component analysis: Application to structural VAR analsysis, *Journal of Econometrics* **196**: 111–126.

Hafner, C. M. and Herwartz, H. (2023). Dynamic score-driven independent component analysis, *Journal of Business & Economic Statistics* **41**: 298–308.

Hafner, C. M., Herwartz, H. and Wang, S. (2025). Statistical identification of independent shocks with kernel-based maximum likelihood estimation and an application to the global crude oil market, *Journal of Business & Economic Statistics* .

Herwartz, H. (2018). Hodges-Lehmann detection of structural shocks – an analysis of macroeconomic dynamics in the euro area, *Oxford Bulletin of Economics and Statistics* **80**: 736–754.

Herwartz, H. and Plödt, M. (2016). The macroeconomic effects of oil price shocks: Evidence from a statistical identification approach, *Journal of International Money and Finance* **61**: 30–44.

Hyvärinen, A., Karhunen, J. and Oja, E. (2001). *Independent Component Analysis*, Wiley, New York.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, *American Economic Review* **99**(3): 1053–1069.

Kilian, L. and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*, Cambridge University Press, Cambridge.

Kilian, L. and Murphy, D. P. (2012). Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market var models, *Journal of the European Economic Association* **10**(5): 1166–1188.

Lanne, M. and Lütkepohl, H. (2010). Structural vector autoregressions with nonnormal residuals, *Journal of Business & Economic Statistics* **28**: 159–168.

Lanne, M., Meitz, M. and Saikkonen, P. (2017). Identification and estimation of non-Gaussian structural vector autoregressions, *Journal of Econometrics* **196**: 288–304.

Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*, Springer-Verlag, Berlin.

Lütkepohl, H. and Netšunajev, A. (2017). Structural vector autoregressions with heteroskedasticy: A review of different volatility models, *Econometrics and Statistics* **1**: 2–18.

Lütkepohl, H. and Schlaak, T. (2022). Heteroscedastic proxy vector autoregressions, *Journal of Business and Economic Statistics* **40**: 1268–1281.

Matteson, D. S. and Tsay, R. S. (2017). Independent component analysis via distance autocovariance, *Journal of the American Statistical Association* **112**: 623–637.

Maxand, S. (2020). Identification of independent structural shocks in the presence of multiple Gaussian components, *Econometrics and Statistics* **16**: 55–68.

Moneta, A., Entner, D., Hoyer, P. O. and Coad, A. (2013). Causal inference by independent component analysis: Theory and applications, *Oxford Bulletin of Economics and Statistics* **75**: 705–730.